Computing Division Self Assessment FY03

1. Division Overview

The Computing Division has carried out an assessment of its performance in several important mission areas. Among the areas of work chosen for self assessment are support of the lab's highest priority and active physics program, Run II of the Fermilab Tevatron Collider. Other areas assessed are the Astrophysics program of the division, the work towards building the software and computing for the CMS experiment to run at the LHC at CERN in 2007, and the support and development work in support of future experiments. We also examined how well we are providing common services and shared resources to the laboratory and to the physics programs.

Mission: The Computing Division's mission is to play a full part in the mission of the laboratory and in particular to proudly develop, innovate, and support excellent and forefront computing solutions and services, recognizing the essential role of cooperation and respect in all interactions between ourselves and with the people and organizations that we work with and serve.

2. Run II Support

2.1 Direct support of the experiments.

The mission of the CDF and D0 departments is to provide extensive computing support for the two Run II experiments, CDF and DØ. Each department includes scientific staff, computer professionals and system administrators to support this mission. The departments have responsibility for the administration and operation of experiment specific hardware such as the analysis clusters, server machines and desktops. Software applications are also produced in these groups and include: database applications, experiment specific data handling, and the GRID tool known as Sequential Access via Metadata (SAM). In addition, members of the departments participate in configuration management issues, experiment user support and in doing production farm processing. The scientific staff contribute to the above activities as well as participating in the scientific life of the experiment by contributing to physics analysis working groups and presenting results at international conferences. The activities of the Run II computing projects was recently evaluated by a FNAL Director's review, and the preliminary close-out found both groups to be very effective in meeting the mission need.

In addition to the direct support through the CDF and D0 departments, the Computing Division also supports the Run II experiment from other departments. The central farms are operated and administered by a highly capable group, and that group also conducts the evaluation of vendors. Database application support and database design, with an emphasis on Oracle, is provided by the division, and are essential services for the experiments. Storage services are also provided by the CD, with development, integration, evaluation and operation services provided. In the close out of the recent director's review, the storage systems dCache and enstore were called out for special congratulations. Network design and installation services are also provided, with high operational efficiency. The FNAL CD also is responsible for computer security for the laboratory, and those policies and decisions affect the Run II experiments. Both experiments rely on central services such as the operators and helpdesk facility and the equipment pool. The high level RunII support goals and assessments follow below.

Goal: Directly support experiments by running central services and systems.

Assessment: The CDF and DØ Departments are highly effective in running the experiment specific analysis systems. Within the rest of the computing division, approximately 15 FTE worth of support effort is supplied to each experiment. That effort is used to supply common services such computing farms, robotic storage, networks, facilities support and customer services. The quality and level of the support is generally very good, and several of the groups supply excellent service.

Goal: Explore commonalities between CDF and DØ to leverage resources and stay in close contact with user requirements.

Assessment: There have been several successful joint projects between the experiments and the CD. These include SAM, Enstore and dCache, and the ZOOM class library and support for software infrastructure for products such as ROOT, C++ compilers, debuggers and development tools. In the coming year we do need to improve the rate of integration of the SAM project between CDF and D0. Next year the CD will also explore ways of combining hardware support as well.

Goal: Oversee the needs assessment and planning for computing resources for the two experiments.

Assessment: The CD implements Directorate reviews of RunII computing at which CDF and DØ are required to submit budget plans for the next year with supporting needs-based documentation. The preparation for and feedback from the review has been quite beneficial to the experiments and the laboratory.

Goal: The Computing Division will contribute to electronics design, development, and testing projects for both experiments.

Assessment: This support effort has been excellent, with contributions to the trigger and luminosity systems for DØ and to the silicon systems for CDF and DØ. These projects are often undertaken to support university groups, and usually involve time pressure as both experiments are in data collection mode.

2.2 Collaboration with the BEAMS Division:

About a year ago the Directorate charged the Computing Division to expand it's mission to more directly aid the Beams Division (BD) to improve delivered RunII luminosity. This additional mission is supported primarily through the "CEPA" department in the Computing Division. Further support is provided from the CCF and CSS departments. Although CEPA predominantly serves the experimental program, expertise in accelerator physics (through SIDAC initiatives), analysis tools, and instrumentation are directly applicable to the Run II luminosity challenge. The high-level goals and assessment of this mission are listed below.

Goal: Analysis of Accelerator Complex Performance.

Assessment: The Shot Data Analysis (http://www-bd.fnal.gov/sda/) project consists of a set of Java-based analysis tools for store analysis. Data from the Sequenced Data Acquisition system as well as from the periodic data sets can be analyzed. For instance, we showed that the antiproton beam lifetime at 150 GeV in the Tevatron was not only affected by beam-beam effects but also by the somewhat restricted dynamical aperture of this large and complex synchrotron. Such conclusions have been reached after numerous analyses, discussions with our colleagues from Accelerator Integration Dept (BD) and from the Tevatron group, and on-going computer simulation done at

Fermilab and Berkeley. Members of the CCF-CD Dept also joined this analysis effort. The SDA system comprises the automatic generation of various summary tables, easy to use data browser and an API for more in-depths analysis of the data from various instruments (Current Wall Monitors, Flying Wires, Beam Position Monitors, Synchrotron Light Detectors, etc.).

We wrote a prototype fitting package for a fully automated Tevatron Betatron Tune Meter. For a beam of relatively small longitudinal emittance the system is able to automatically report betatron tune values to a fraction of one Hz. This tune fitter is partly based on fitting tools developed in HEP. Its current implementation is based on the ROOT software. For more details, see (http://www-bd.fnal.gov/tevtune). This new instrument has recently been used to provide accurate and automated measurement of the chromaticity and of the beta function at selected locations. As usual, this work has been done in close collaboration with the Beams Division staff. We plan to rebuild the system to provide enhanced maintenance and better performance.

We designed, developed and deployed a snapshot monitor for aiding in monitoring and diagnosing the state of the Booster Accelerator. The tool is working well and is currently in use. While developed specifically for the Booster, this tool was designed from the beginning with the idea it could be useful in other accelerators throughout the laboratory.

In summary these analysis tools were effectively used to advance the understanding of the Accelerator Complex.

Goal: Modeling of the Accelerator complex:

Assessment: CEPA staff that are members of the SciDAC funded Advanced Accelerator Modeling team, have worked on developing Synergia, a 3D parallel code suite for beam dynamics simulations. These tools were used to model the FNAL Booster, with the objective to understand Booster losses and study approaches to minimize them. This work has been recognized as essential for the success of the FNAL program by both the DOE and the Fermilab management. During the last year development of the space-charge code needed to model collective beam effects early in the Booster cycle was completed. Experiments at the Booster to compare these models with data were performed. In the process of making these measurements, we developed and implemented calibration technique for the Booster's Ionization Profile Monitor. Our results were presented at the ICAP and PAC conferences; we have submitted a paper at PRSTAB. Overall, our effort was successful both in developing and in applying our simulation tools to the Booster. We also ported or code to both commodity PC clusters and specialized parallel machines in addition to developing a build system and web based documentation, thus providing the infrastructure for a broader use of our tools.

In summary this suite of modeling tools made important contributions to the understanding of the accelerator complex.

Goal: Automation and Upgrade of Measurement and Control Systems:

Assessment: We also participated in an investigation of the BD Vax Migration effort and evaluation of the related new Java Based Data Acquisition subsystem. For instance, we wrote a small D.A. application targeted for specific Tevatron studies, where the experiment results can be cleanly expressed as a set of tables of actuators values and corresponding beam property measurements. This work is solely based on this new Java-based D.A. In the course of developing the SDA system, the Tune Fitter and other applications, the performance of the new D.A. has been studied. We have had valuable discussions with members of the Controls Dept on possible improvements.

We wrote a data acquisition system for the Flying Wire instrument for the Recycler. However, as we made good progress on this LabView based software, we learned that the first installation of the instrument hardware in the recycler during the January shutdown failed because of vacuum problems. Final resolution of this problem, via a thorough vacuum certification process, is still pending. Unfortunately, these setbacks have cast a shadow on the ultimate usefulness of this work. It is hoped that this work will serve as a base for an update to other flying wire installations in the other machines.

This year (FY03) the Beam Position Monitor system for the Recycler Ring was installed and commissioned. The CEPA and CSS departments played a major role in the design through commissioning phases of the project, and particularly the calibration system. This project was successfully executed in FY03.

Late in FY03 the Computing Division took on the major challenge of leading the effort to upgrade the Tevatron Beam Position Monitor system. External reviews have noted this upgrade as critical to future success of the Tevatron. The resources to realize this system primarily reside in the CEPA department, with important contributions throughout the Division. Work in FY03 focused on generating detailed requirements. Progress on this project in FY03 was adequate.

In summary good progress has been made in advancing the instrumentation of the accelerator complex, with the exception of the flying-wire readout system which has yet to be deployed. Once the detector issues are resolved in FY04 the flying-wire readout system will then be implemented.

3. LHC support – CMS

3.1 Overview

The CMS department of the Computing Division is central to support for the US-CMS Software and Computing Project. The department is involved with all aspects of computing from the deployment of basic hardware components and system services, to the integration of those facilities into a grid-enabled distributed computing infrastructure, and finally to developing the applications that run in the environment. The activities of the department can be divided into four equally important activities: developing CMS core software, providing computing resources to US and international

CMS, developing computing infrastructure and supporting the US Tier-2 centers, and developing and supporting analysis activities for physics. We will briefly describe this year's progress and challenges in each of them.

Goal: Contribute to the development of CMS core software

Assessment: Fermilab has made a significant contribution to the development of CMS core software, due mainly to the quality of the people hired. This year a Fermilab developer was responsible for changing persistency mechanisms used by the CMS core software framework, COBRA. This was one of the most important software tasks of the year. It is also central to the operations of international CMS and relies heavily on packages developed at CERN. Fermilab developers were also responsible for developing the tools that provide binary distribution of CMS software for all simulated event production and the development of the tools used to specify and submit production jobs at all processing sites. Despite the geographic separation, Fermilab-based software developers have remained well integrated in the software development effort.

US CMS has a funded but unfilled software engineering position at Fermilab. This position needs to be filled or a new person hired as soon as possible. There are several important, high profile tasks that US CMS would be capable of performing, if the position were filled. CMS has a need for software development to modify the CMS core software to perform more efficiently in a grid environment, including application level monitoring, resource prediction, and check pointing. Additionally, CMS needs to develop expertise to interface the CMS metadata to developing distributed analysis environments. These tasks are a natural fit to Fermilab due to our existing activities in core framework and production and knowledge systems.

Goal: Provide computing resources to the CMS collaboration

Assessment: The second responsibility of the CMS department is to meet US CMS obligations to the experiment by providing computing resources. Those expectations are generally met. We performed a large procurement of processors, storage devices, and network switches this year to bring the Tier1 facility in line with CMS expectations for computing this year. In an ideal world we would have had the resources commissioned before the start of simulated event production, which began in August. However, due to new personnel coming up to speed and the length of the procurement process we expect systems on the floor by October.

The new procurements significantly increase the number of systems to be managed. The Rocks configuration management software, which was deployed last fall, helps with operations, but does not eliminate the need for operational support. We have been operating with 1.5FTE of operations support, but this has been a struggle. With the additional systems we estimate we need 2 FTE for operations this year.

	2002	2003
CPU	75	200
Storage	10 TB	34 TB
Throughput	200 MB/s	700 MB/s

Goal: Develop the grid computing infrastructure necessary for support of Tier 2 centers in the US working with the Tier 1 center at Fermilab

Assessment: The grid computing infrastructure development has progressed well. We have deployed a basic grid enabled system for generating simulated production at the prototype Tier1 and Tier2 centers. We have reduced the effort required to run production. We have a project to extend the functionality of the system to manage Virtual Organizations (VO) on the grid, and we are deploying a multi-VO environment called Grid2003 in coordination with US ATLAS, SDSS, and LIGO.

While there has been interesting progress, it is difficult to maintain with the current manpower. Again, the US CMS project has budgeted additional positions, which have not been realized so far. The need to succeed with Grid2003 has introduced a number of external constraints, which are further stressing the available people.

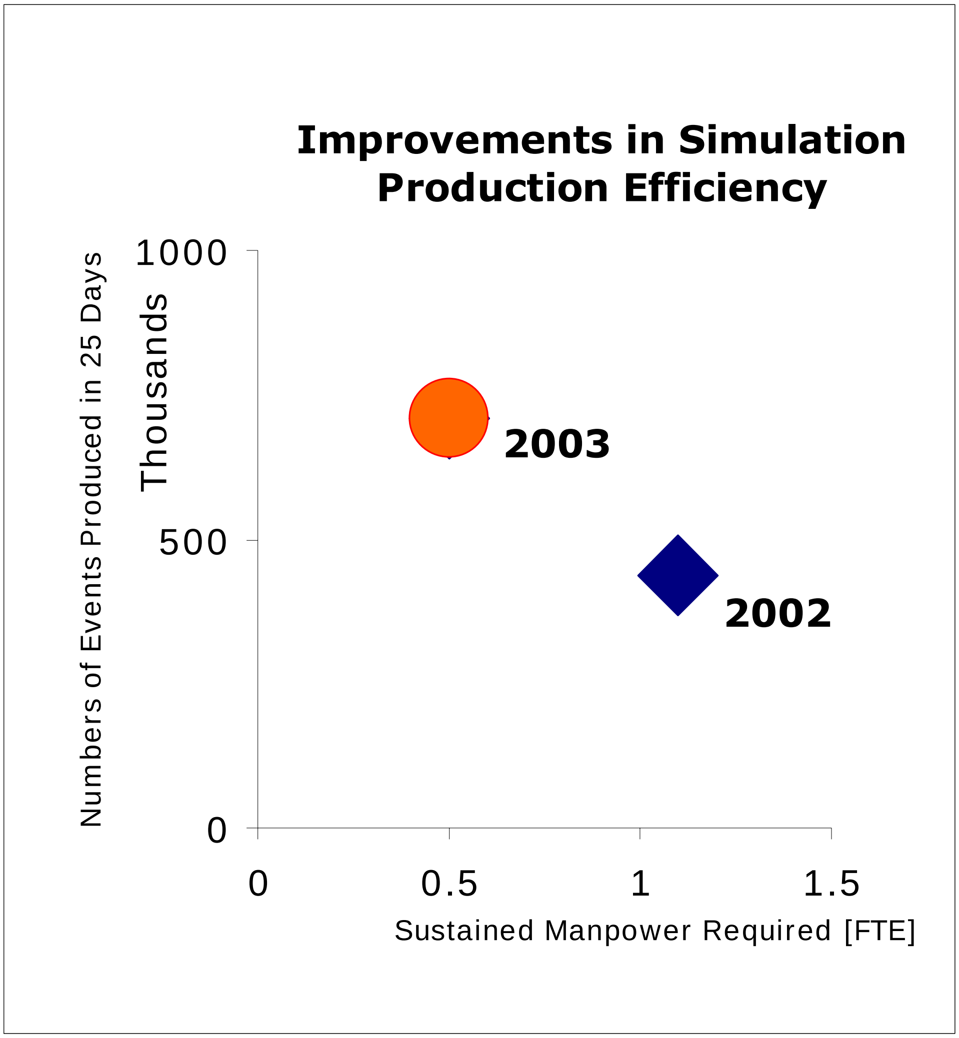

Goal: Develop and support analysis capabilities in the U.S.

Assessment: We are providing and developing services to facilitate analysis. The current farm and storage capacity are small, but improving with new procurements. The number of active users is also small. We've been working to determine if there are not a lot of analysis users currently working in CMS, if we haven't advertised sufficiently, or if we are not providing attractive services.

We believe that our analysis facilities are becoming rather advanced. The interface to mass storage over dCache is transparent and works well with the CMS core software; there is good user access to temporary storage and the new procurements will provide large, reliably, and long lived user space; and there is sufficient user processing capacity including access to (legacy) Objectivity data. Until recently the CMS software releases were difficult to keep synchronized with CERN, due to the frequency of releases. A technique was recently employed to allow the CERN distributions to be visible and usable at the Tier1 over afs. This should improve the situation considerably by giving transparent access to all releases and pre-releases as soon as they are published at CERN.

We were not able to provide test beam users with the level of service we had hoped this summer. Some of the test beam computing requirements were small and could be satisfied with local desktop support, leading users not to make requests for central support. Also, we did not have sufficient effort to provide all the services we were interested in. These small time critical analysis activities, like test beams, should be a very good test of our analysis services. We hope to have more success in the future.

We expect the number of analysis users will increase significantly as the physics groups prepare for the technical design report. In order to offer reasonable services the CMS group is going to need to devote more effort to user analysis. While this is one of the primary goals of the CMS Tier1, it will further stress our available people.

4. Astrophysics support

The mission of the Experimental Astrophysics Group within the Computing Division participates in particle astrophysics experiments that require a significant computing capability. Currently the group works on E885 (Sloan Digital Sky Survey), which is entering its 4th year out of a planned 5 years of operation. The experimental facilities are located at the Apache Point Observatory (APO) in New Mexico, while data processing is all done at Fermilab.

Goal: Continue to support the data acquisition system, software development, survey planning, data processing, and data distribution to the SDSS collaboration and to the public at large.

Assessment: The group receives some support from CCF for assistance with the data acquisition system and from CSS for system support of computing machines located in FCC. Much of the data analysis software is provided by other institutions in the SDSS collaboration. Data processing and handling make extensive use of the fixed target farms and the Enstore tape storage systems. Certain aspects of data processing are time critical because of the need to support operations at APO; in particular, imaging data must sometimes be processed and used to design "plug plates" on a month time-scale.

The data acquisition system is in a maintenance mode, with only occasional bug fixes and small feature enhancements being provided. The system is old and not entirely reliable, but at the moment it is not a major source of lost time. Concerns remain for the longer term future, however.

The essential aspects of data processing are working well. The median time to turn around imaging data has been reduced from 20 days a year ago to about 5 days this year. Other smaller datasets are also being turned around well within the needed timescale of the experiment. These gains have been achieved by making processing highest priority, by focusing on improving the data processing infrastructure, and by benefiting from a "Moore's law" increase in processing power, disk storage capacity, and network speed. A complete reprocessing of all data has recently been completed as well. The areas of quality assurance diagnostics and regression testing of software code are less satisfactory since their development has had lower priority. These are expected to be addressed in the coming year.

With the advent of inexpensive terabyte-sized file servers and the maturing of the Enstore tape robot systems, data management has become much easier. All processed data, including an average of two reprocessings, are kept online and are available to the collaboration. Raw data tapes are copied into the tape robot as they are delivered from APO and are also available to the collaboration with a bit of effort. Older raw data tapes from early in the experiment are not yet in the robot, and accessing them is difficult, although they are seldom used.

Data distribution has two components: flat files and a database. The database (called the Catalog Archive Server [CAS]) contains object catalogs that are also available in the flat files but organized for efficient queries. To first order data distribution is successful: all files are available to the collaboration, and subsets of the files are available via web and rsync interfaces to the public. (It should be noted that unlike particle physics data, astronomical data have utility and legacy value to the astronomical community at large.) There is a web interface with good documentation. The CAS has been a mixed success. It is being developed by another SDSS institution and has suffered from feature creep and a switch in database technology. At present there is a CAS version that contains the subset of data catalogs that have been released to the public, but not the full processed data. In fact, for the most recent public data release, an independent CAS with limited features was developed at Fermilab to provide backend support for the public data web site. To improve the situation, a more formal project management approach is being used to monitor the database development and deployment at Fermilab.

5. Future experiments support

5.1 Overview

This area covers the computing support for the Fermilab experiments other than CDF, D0, CMS and SDSS. There are currently four projects in this category, MINOS, MiniBOONE, BTeV and CKM. These four experiments are supported out of the Experimental Support Department.

Goal: Provide support for computing of the MINOS experiment

Assessment: The construction of the MINOS Far Detector at Soudan, Minnesota was completed in June 2003 and the magnet was energized on the second supermodule. The experiment is now taking physics quality data with the Far Detector and data analysis of cosmic ray data from is underway. MINOS is using the central computing resources at Fermilab, namely the STKEN tape robot, STKEN DCache, FNALU/AFS cluster and the “Fixed Target” production farm. Support of these facilities is provided by the CSS and CCF departments and they have responded promptly to the experiment’s requests over the last year. The department provides direct support for MINOS computing with one scientist and three computing professionals. They are responsible for user consulting, some of the offline software development, data handling, Control Room Logbook support, Linux support and equipment purchases to name a few. There were some personnel changes in this group in the last year, which have resulted in a better distribution of the necessary skills needed for MINOS support. One of the main focuses of MINOS support in the next year will be design, procurement and deployment of an analysis system and deployment of the SAM (Sequential Access with Metadata) data handling system in use at CDF and D0.

Goal: Participate in and provide support for the MiniBOONE experiment.

Assessment: The MiniBOONE experiment has one scientist in the department. They are using some of the central facilities, namely the STKEN tape robot and AFS but also have their own Linux cluster that is supported by the experiment.

Goal: Provide support for research and development activities aimed at preparing the proposal for the BTeV experiment for external review.

Assessment: The direct support for BTeV computing consists of one scientist and one computing professional. This is currently a fairly modest effort that will clearly need to ramp up significantly once the experiment is approved. The group has been working on various projects included the design of the interface between the offline software and the Real Time Embedded Systems (RTES) project with members of the CEPA department, which will control the trigger/offline farm on which the offline software will run. The reuse of the trigger farm for offline processing is a novel feature of BTeV. They have worked on Tutorials to educate BTeV physicists on the use of the software suite and to smooth the transition to C++ for those who have yet to make it. The group has also been supporting the test beam effort for the EM calorimeter at Protvino.

Goal: Provide a small amount of support for activities related to potential future experiments at Fermilab, such as CKM and Off-axis neutrino experiments.

Assessment: has one scientist in the department. Like BTeV their support needs will ramp up once the experiment is approved. They are also making use of the STKEN tape robot. They also are working on development of the packaging of the TDCs and ADCs on PCI cards in conjunction with the CEPA department.

6. Computing and Engineering for Physics Applications

The Computing and Engineering for Physics Application (CEPA) department has a diverse mission and has been charged to create and support core components for physics applications. These components can be software, both online and offline, or hardware (electronics). We serve the experiments and have been recently to work with the Beams Division.

This assessment document is divided into parts that align with the sections inside the department. At the end there is an assessment of several activities that cross the boundaries including the work with the Beams Division and the work on Outreach.

6.1 Engineering and Engineering Support

Goal: Support ongoing HEP experiments by providing diagnostic, repair, and upgrade support. Recently the section has been called on to design and produce instrumentation for the Tevatron and the Recycler.

The group has provided supported for a large number of projects this year including the following:

- At CDF: Silicon Vertex Detector DAQ and CDF DAQ Event Builder system. The experiment is efficiently collecting data and the hardware is operating reliably.
- At D0: Trigger Distribution and Serial Control System, Central Fiber Tracker (CFT) Mixer system, VRB Controller module, and Level 2 Trigger Alpha and Beta processors. - The experiment is collecting data and the above hardware is operating reliably with the exception of the VRBC. The VRBC has been orphaned by the previous designer and lacks the documentation that would facilitate problem diagnosis.
- At the Fixed Target lines: E907 engineering consultation and the BTeV pixel tests at the Meson Test Beam Facility. For the BTeV effort the hardware is tested and awaiting beam.

Assessment: It was overall a successful and productive year. Support for the Run II experiments was a primary focus. However, there were a few setbacks year. Two engineers from the group moved to positions elsewhere and it has not yet been possible to replace them. The number of projects has not decreased, and the rest of the group has responded well to fill in the gaps.

The ESE section contributed to several upgrades projects at the collider experiments and the Accelerator complex including CDF DAQ, D0 L2 trigger, D0 L1 Luminosity firmware and the Recycler and Tevatron Beam Position Monitor (BPM) upgrades. Most of these projects are still active and making acceptable progress. The hardware for Recycler BPM project was completed and installed this summer and is being tested as the accelerator schedule permits. The work for this project was completed on schedule. The Tevatron BPM project is just getting started and involves collaboration with Beams division that is going well.

6.2 Applied Physics Software

Goal: To develop and support a body of common HEP library software, including ZOOM and CLHEP; to assist Run II and other experiments with offline software development and to support the experiments in the areas of database and information systems related software.

Assessment: We have developed a strong team of experienced C++ experts in the Software Library Development (SLD) whose expertise is in high demand at the experiments. They are responsible for a body of library software, including ZOOM and CLHEP. While this was originally as a responsibility to Run II, we are presently moving the focus substantially toward making this software more valuable to other HEP users such as CMS as well. We are in the completion stage of a repackaging aimed at the world of scientists who choose not to depend on SRT. Because this inherently requires satisfying a global spectrum of preferences, we found this task technically challenging with no single "best" solution; we have succeeded in evolving technically a sound solution that meets all the vital criteria. Related to this is the job of providing C++ expertise, advice, and focus for FNAL developers. A very important component of that is our participation in the J16 C++ Standards committee, which is beginning to pay off for the entire HEP community. We have not only introduced improvements which will be important to the scientific community, but also have reversed a trend of the C++ committee de-emphasizing the needs of the scientific community. The reason for this success is the hard work and high motivation of two FNAL representatives.

Goal: Develop and support database and information systems related software.

Assessment: The use of databases is expanding at the laboratory and the group could use at least one more database design expert to keep up with the growing demands.

Database applications include many aspects of data management calibration, trigger, luminosity, and configuration database development and support. The tasks include building and maintaining middle tier servers including those for SAM and general-purpose read-only database information. SAM is the data management system being employed for D0 and CDF or Run II. The group has several additional responsibilities in the SAM project, in particular dealing with the user interfaces and some operational aspects. The degree of success of SAM support is indicated by the fact that CDF has committed to use this software and the group has made modifications needed to satisfy CDF's requirements.

In addition the group works closely with the experiments to provide documentation and code for client applications that use the database information. Work has largely centered on Oracle applications, but the goal is to provide general interfaces that allow for multiple commercial and freeware database solutions as well. The group also provides a monitoring framework and tools that allow the database delivery systems to be alarmed, monitored for performance, and diagnosed for problems. This set of tools came about because of realization by several people across the whole section that multiple efforts to improve monitoring could benefit from a reasoned, coherent shared effort. The monitoring tools came to fruition in less than a year, and are paying dividends. For example, there is now daily identification of each experiment's most heavy users; by working with them to reduce inefficiencies, we have brought database crises under control.

Recently, the group has been helping to review parameter databases being used in the beams division to track magnet and other information for Tevatron operation. This appears to be a project that will continue for the next several weeks and we hope will play an important role in improving the operation of the accelerator. We are currently making recommendations concerning moving some of their existing data to a new, more maintainable database, unifying into one modern database product such as SYbase, PostgreSQL, or Oracle. This database would be prepared for the influx of new survey data starting in mid-fall.

6.3 Physics Simulation and Modeling

Goal: To support the experiments, the accelerator and the theory group by providing support for tools used in the simulation of physics events, in the simulation of detector effects, for the simulation of beams effects and for calculations on the lattice.

Assessment: This group is not large, but is able to successfully make contributions in a wide range of activities. The group would be stronger given more resources. The group acquired additional expertise on the use of the Geant4 simulation package through work on. Developing a full Geant4/OSCAR simulation of the CMS hadronic calorimeter test-beam experiment. As one of the main contributors to the Geant4/OSCAR validation effort, the group was able to make an importance contribution both to the CMS physics analysis and to the Geant4 development and make the Geant4 library available to the FNAL users. In addition, work continued to maintain the Geant4 based beam simulation tools developed by our group for Ionization Cooling Modeling. This support allowed the MICE and MUCOOL collaborations to continue their design work using our tools.

The group is responsible for maintaining and testing the physics generators and related utilities needed for collider physics analysis, for Run II and beyond. The expertise provided by the group is a valuable asset to the experiments and helps to ensure the fidelity of the Run II physics simulations.

A series of meetings with the experiments were organized and as a result a database of "blessed" Monte Carlo sets together with their generating parameters was implemented. This effort could be strengthened with some additional manpower, perhaps from guest scientists or a post-doc. Member of the group were primary contributors to the international effort to systematize comparisons of different physics generators and acted as consultants for related issues for LHC study groups. The focus on generators in the computing group at Fermilab is somewhat unique and we feel essential.

The group provides well-tested libraries of common generators for distribution in the FNAL environment. The libraries are supported through mailing lists and members of the group act as the interface to the original authors. The users (including Run II experiments) appear satisfied.

The group also plays a leading role in providing infrastructure and support for the FNAL Lattice QCD facility as well as acts as one of the main users of the facility. The resulting simulations of the Charmonium states produced precision results that were presented at the Lattice and Lepton Photon conferences, and were submitted to PRL.

The group is also a member of the SciDAC funded Advanced Accelerator Modeling team. Through this collaboration, they have developed Synergia, a 3D parallel code suite for beam dynamics simulations. This code has been successfully used in the evaluation of space charge in the Booster and will provide a most useful tool for future accelerator studies. The group, though small, has been very successful and productive. This is the area that would profit most from more resources so that the accelerator modeling can proceed in a timely fashion for Run II.

6.4 Online and Analysis Applications

Goal: To develop and support real-time and analysis software for the experimental community and for projects oriented towards improving the Tevatron's performance.

Assessment: Our responsibilities to the Run II experiments have proceeded well. We've been heavily involved in performance monitoring and integration efforts on the online system at D0. This proactive approach has helped to minimize downtime over problems that spring from hardware failures and the occasional problem with software upgrades. We successfully deployed an enhancement request to increase the number of L3 trigger bits. We have been comfortable enough with the software to move the support load to a more junior member of the department, freeing up expertise to work on other projects. Needed changes to CEPA developed software for CDF have been minimal and timely, and we have added some additional responsibilities in building front-end kernels specific to CDF. The front-end support at both experiments (CAMAC, VxWorks) has been extremely stable.

For non-Run II experiments, we worked on these projects as resources were available. The work that we've provided here is well received and appreciated, but we are behind on some of our goals because of lack of resources. Over the past year, we have been actively involved in the BTeV approval process, making sure all necessary project scheduling and documentation is available for upcoming reviews. We spent time in evaluating software that could be used in the upcoming DA system and some effort in defining the requirements, as people-time is available. For CMS, we have (co)developed and supported the data acquisition system used by HCAL and EMCAL with the end goal of a successful HCAL test beam effort this summer at CERN. We helped E907 develop and deploy a data acquisition system that is currently in use at the experiment. We have filled in gaps for experiments in the test beam effort including providing a test beam facility DAQ (in progress) and OS support of the DA machines.

We continued to support the electronic logbook (CRL) for D0, MinibooNE, and Minos. We deployed it at CMS this year as well. Two major criticisms of the CRL logbook were addressed this year by improving the installation procedure and a web based interface for adding new entries that will be deployed next month.

Our support of the physics analysis software tools has been very successful. The Run II experiments have elected to rely on tools that are widely used in the physics community but created at other labs. By providing a strong local support for those tools, we have been able to achieve a very high turn-around rate on bug fixes and we also have been able to insure that features that were needed by the Run II experiments were implemented in a timely fashion.

We also have been providing support for 3D data visualization. In particular, we have been collaborating with MiniBooNE to develop their Event Display. This work has been somewhat delayed due to the load of planning and preparation for SC2003.

7. Core Support Services

The Core Support Services (CSS) Department within the Computing Division has many major missions for the Laboratory and Division. It is involved in operations of almost all services the Division provides and provides many core services and support functions directly. These responsibilities include:

- Base support and vendor liaison for all the supported operating systems;
- Global infrastructure support for the supported operating systems (e.g. Windows domain controllers, license servers, boot servers, etc.);
- Lab-wide core services (e.g. email gateways and IMAP servers, printer service, AFS global file service, software product distribution, etc.);
- Centrally-managed web servers, including the main site web server;
- Shared access computing facilities (batch and interactive) for scientific and technical analysis and development (e.g. FNALU, etc.);
- Dedicated computing facilities for production analysis of physics data (e.g. farms);
- Standardized software applications for distribution to FNAL machines and offsite collaborators;
- Recommended configurations for secure and reliable system operation;
- UNIX and Windows workgroup support and administration as described by negotiated MOU with client organizations within FNAL;
- Computing support for conferences sponsored by and/or at FNAL;
- Consulting services on system configuration, hardware procurement, computer security, and system administration training/hiring;
- Operation of the Computing Division’s HelpDesk, including development, support and administration of HelpDesk applications and databases, and tools to assist in the automation of the Division’s operations and monitoring.
- Development and support of database design and management tools;
- Development, support and administration of bookkeeping databases and catalogs for experiments (CDF, DZero, CMS, Auger, MINOS) and other laboratory clients using Oracle and open source MySQL or PostgreSQL databases;
- Development and support for the Computing Division MISCOMP information management systems and other Division infrastructure;
- Repair, diagnostic, calibration and advisory services on a wide range of specialized high-energy physics instrumentation, electronic modules, computing equipment and computing peripherals, especially when vendor repair is not available or is limited or costly;
- Management and tracking of over 40,000 electronic modules and more than 20,000 individual pieces of computing equipment in the Physics Research Equipment Pool (PREP) equipment pool;

- Administration of Lab-wide maintenance contracts for computing hardware and software and tracking of software licenses through the equipment database;
- Administration of the FermiTools software technology transfer program.

Goal: Provide efficient and reliable core computing and equipment services

Assessment: The Department continually monitors and assesses performance in these functional areas, based on both client feedback and "hard" metrics, which are used in our on-going program to make improvements. Organizational and process changes, and new technology, have been used where appropriate to accomplish these improvements. The new organization of the Department has greatly contributed to the ability to flexibly bring together resources from anywhere in the Department to address problems or undertake new projects.

A measurable success has been the Department's contributions to the success of the Run II experiment's data-taking and analysis in the areas of database administration and production farms processing. Another success has been our contributions to the Run II luminosity upgrades through support for the hardware modifications for the beam position monitors.

One current area of focus is reliability and maintainability of commodity hardware for large computing clusters, where we continue to improve our processes for evaluation of hardware and certification of vendors. We are also working with vendors to improve overall uptime through more robust hardware and software and lower time-to-repair.

8. Computation and Communication Fabrics

8.1 Overview

Computing for the physics programs of the laboratory as well as for the technical and administrative work of the laboratory relies on certain core capabilities that are provided for everyone. This computing fabric provides a campus-wide network and many computational facilities.

8.2 Computer Security Team (CST)

The Computer Security team provides a dedicated technical organization to support the Computer Security Executive, and assist line management in maintaining computer security at Fermilab. This is done under the governance of the Fermilab Computer Security Program Plan (CSPP). The CSPP articulates the Lab's Approach to Computer Security. The CST provides (among other items) training of System Administrators, a strong authentication infrastructure providing Kerberos 5 and X.509, network-based scanning and monitoring of network activity, a framework for incident response, a framework for identifying critical vulnerabilities and determining the labs response in mitigating them.

The Computer Security Team also tracks and contributes to computer security technology developments, and attempts to play an appropriate role coordinating computer security among significant HEP facilities, so as to reduce the impact of various security measures on the collaborative and increasingly inter-operating world wide HEP collaborations.

Goal: Provide a conceptual framework for the lab's computer security activities, based on the Fermilab Computer Security Program Plan (CSPP)

Assessment: A rich framework exists, and is very well grounded in the prose and spirit of the Fermilab CSPP.

Goal: Provide central technical infrastructure as apropos, for example, a Kerberos key distribution center (KDC).

Assessment: An appropriate technical infrastructure exists. Because most of the work is "collaboratively open", and because the nature of the lab's scientific program requires extensive collaboration with many other institutions, a significant amount of this infrastructure is realized as scripts built upon publicly available tools. The apparent path forward has the CST depending on a significant amount of in-house software. The Team needs a plan and direction for this effort. The CCF Department is seeking to augment the group for a person with suitable potential to address this issue.

Goal: Provide business process apropos to a central group (e.g. exception processing, training, awareness program)

Assessment: Business processes exist, and were reviewed in the self assessment. It has been noted that the group has been suffering from a lack of available manpower, with an impact in turn around, and limiting the expansion of process. This is being investigated.

Goal: Bi-annually conduct a peer review and self assessment.

Assessment: This has been done, with the results incorporated into the Team's work plan.

Goal: Understand and steer relevant technologies where appropriate.

Assessment: Staff of the team are working in several areas, including participation in the Global Grid Forum and involvement in software engineering, contributions to the MIT Kerberos Distribution, and participating in various inter-lab fora.

8.3 Networks

The Computing Division supplies the scientific program with networking for its Data Intensive Science, and provides general networks for a considerable fraction of the site.

Because of the collaborative, multi-institutional nature of High Energy Physics, special attention is paid to off-site connectivity, with a focus on the remote-site data needs of experimenters and theorists, supporting a large number of visitors with laptops (inducing the need for good support for registration, network analysis and incident response).

Ethernet is the basis for almost all of the lab's very considerable scientific computing infrastructure. The design of data-intensive networks has been and remains a very important contribution to the lab's physical infrastructure.

Goal: Furnish a centrally-managed campus network, continuously evolved to meet modern usability expectations, having a configuration allowing for central management, and low operational costs.

Assessment: Such a high-quality managed network exists, and has the requisite properties. Challenges of adequate management were met in relevant areas, including management software enforcing elements of proper use (autoblocker), increased Netflow support , migration of monitoring to HP OpenView, tracking core use and advocating 10 GB modules as apropos, and extensions of the inter-building fiber plant. The Core Network Computer Security Plan was revised.

Goal: Furnish a designed network for the lab's data-intensive computing, including Network-based storage systems, network-based analysis clusters, and production farms.

Assessment: This has been an outstanding success. Such networks are now deployed or planned to be deployed in every appropriate area.

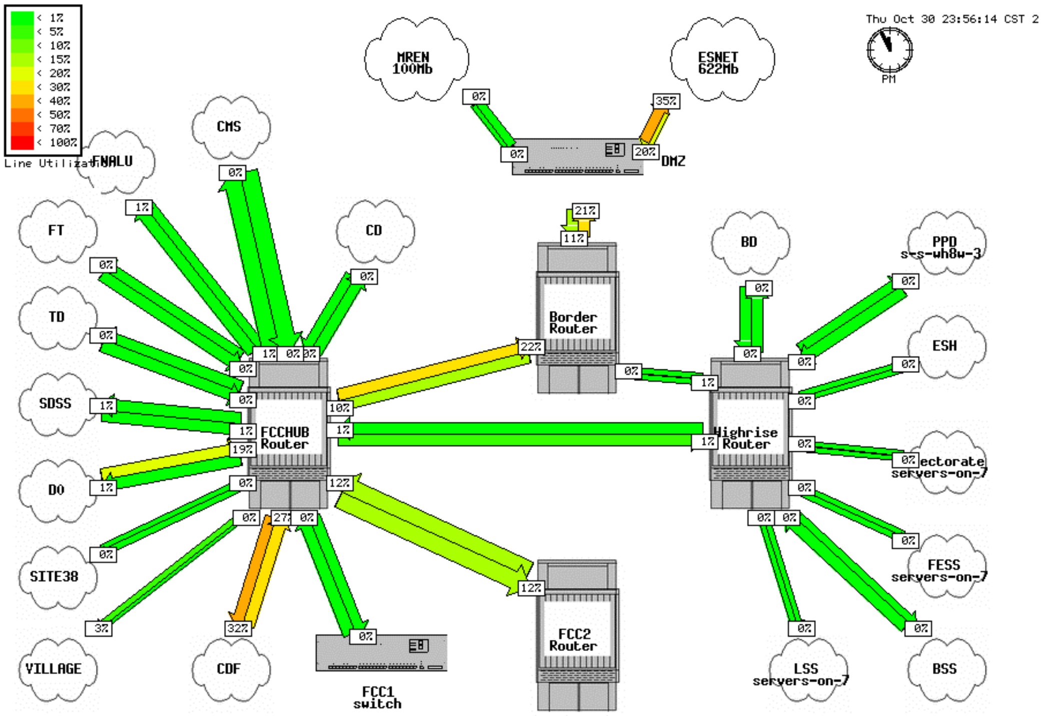

Goal: Furnish off-site connectivity apropos for the production needs of the lab's current experiments, and for inter-lab systems development and demonstration needs.

Assessment: Production traffic is to be handled by ESNet. An OC-12 upgrade was obtained in time to support the lab's production traffic. A Fiber RTU to the Starlight Optical interchange facility is in process, but is proceeding more slowly than hoped for. This RTU will allow the Lab to connect to research and production networks with novel technologies and at very high bandwidths.

8.4 Storage and Data Movement Facilities and Infrastructure.

The CCF department builds and operates facilities for permanent automated storage of designated scientific data sets and providing for access to the data with appropriate performance. Currently, these facilities are implemented as automated tape libraries operated by the department, and disk caching software run on (typically) experiment-furnished and administered disk-full computers. The tape systems hold over 1 Petabyte of data, with more than 10 TB of data being transferred to or from tape each day. Experiments implement approximately 100 TB of disk. The disk systems typically transfer over 20 TB/day to a more chaotic physics analysis system. All data move over an Ethernet data Network.

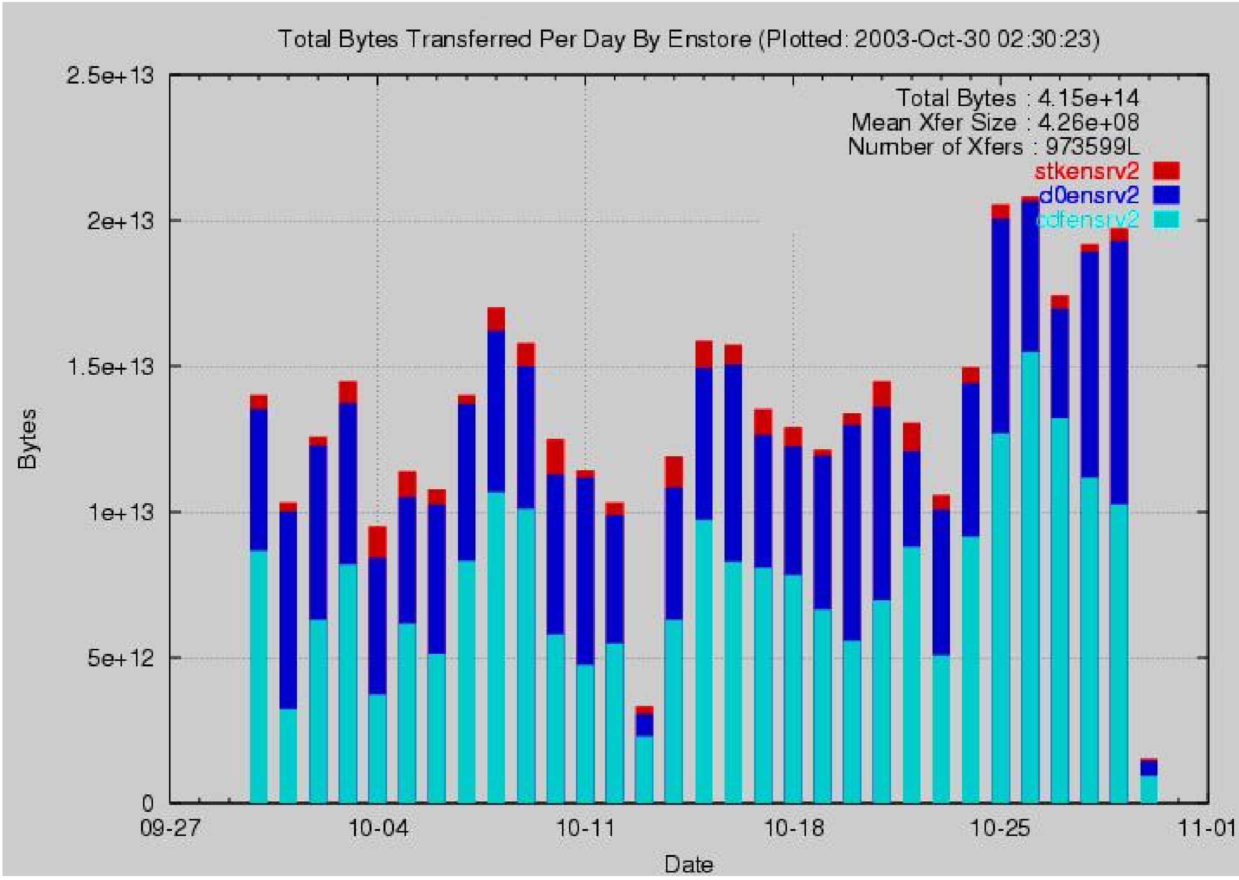

Goal: Maintain capacity and functionality to satisfy the FNAL experimental program.

Assessment: Capacity and throughput have met experiment's needs very well. This was confirmed in the Director's review of Run II computing. Success of the FY03 program has engendered more desiderata for FY04.

Goal: Maintain the operational characteristics of the permanent store (tape system).

Assessment: The tape system has very good availability in the face of a very large increase in capacity, the introduction of a new tape drive technology, and the expansion of the systems to ~70 tape drives. The state of the system, including many operational parameters is monitored continuously by a "primary" system administrator and reviewed in a meeting that includes storage system administrators, software developers, and department management.

However, experience from running the deployed systems and external requirements have changed and extended the notion of a solid operational foundation. The Department is investigating the following items and regards them as important in FY04.

a) Increased verification of the integrity of tape meta data backups.

b) Moving from an "iterated install" method to a boot image method.

c) Examining the physical computer plant for unwarranted diversity.

d) Fully implement using the tape write-protect tab for reasons of safety and security.

Goal: Expand the designated data sets supported by the systems to legacy data set with substantial active interest.

Assessment: This requirement has emerged late in FY2003. A technical proposal exists and work is under way.

8.5 Lattice Gauge Computational Facility.

The FNAL Lattice theory program compliments the FNAL experimental program. A basic tool of lattice theorists is simulation. Practical simulations require coherent calculations over many processors, with a requirement for low latency communications. This is unlike the requirements for the FNAL experimental program, where calculations are more farmed out to individual computers and can proceed independently.

Goal: Provide a facility adequate for the FNAL theory program, and appropriate to the lab's role in the SciDAC lattice collaboration.

Assessment: CCF has outstanding expertise in the area of clustered computing. It has extended the lattice computational facility. Anticipating use in the context of the SciDAC collaboration, it has begun collaborative activities to form a data grid to handle Lattice data sets.

8.6 Project Capability

The department's project capability is substantially located in two groups, the Middleware Applications (MAP) Group, and the Fabric Technology Projects (FTP) group. These groups participate in projects associated with the groups and facilities mentioned in this section of the self-assessment. The groups' staff work in the context of the Computing Division Matrixed projects, where there expertise is salient, or where work and expertise gained in the project helps inform the department's own technical program, or simply where their is a need.

To illustrate with examples, the department works in the area of Virtual Organization (VOX), Grid site security (SAZ), Grid Job Scheduling (JIM), systems concepts underlying data center facilities expansion (blade computing investigations), and assistance to the Beams division (Debugging SDA, Emittance Class)

Goal: Effective project formulation, including effective check points, mitigating the risk of perpetual involvement, while making sure that projects are delivered and put into production.

Assessment: Very good motion: The MAP group leader has been able to insert herself to assist the formal project manager when indicated (SAM JIM). Long running projects, with fuzzy deliverables are now thought to be a program of projects, with projects held to short term milestones tied to a production goal.

Goal: Development of the staff, with an emphasis on technologies and soft skills important to collaborative work. Ensuring a project environment apropos for CCF staff.

Assessment: Very good motion in a year. Staff that had worked in in-house Python projects now work in Java, and web services which seems preferred for grid collaboration and beams work.

8.7 Operations

Operations refer to providing computer operators for the computers and peripheral devices located in the Feynman Computing Center. Computer operators also do ancillary work, due to their familiarity with the data center and their presence in the center beyond working hours.

This group was moved into the CCF department in August and we are still evaluating its goals and assessing its performance.

9. Division and facility infrastructure

9.1 Overview

The Computing Division underwent a large reorganization and consolidation in early FY03.

Goal: To consolidate services and to leverage similar expertise for broader applicability across projects.

Assessment: The number of departments was substantially reduced and people were moved into these departments according to functional affiliations. The results of these moves have been very encouraging. Personnel allocation to projects has become easier and has allowed flexibility in adjustment. There are fewer points of contact for operational issues and communication flow is more efficient. We believe that we are on the right track and are now in fine tuning mode.

Goal: Provide oversight, recommendations and tracking for various division wide operational activities. Manage office and computer space, building utilities and safety programs.

Assessment: Personnel relocations to different offices were handled efficiently and with minimal interruption to work.

In FY03 we provided the necessary physical infrastructure for the addition of 750 computers for CDF, D0 and CMS experiment computing needs. We are now in the process of planning computer room space for the next several years. We have introduced the first phase of computer room automation which will allow unmanned operation for at least 1 non-prime shift per day. We are in the planning phase for a web-based facility management system for computer room equipment, space and power management.

In FY03 we achieved greater than 95% of ES&H training throughout the division. We have averaged two departmental ES&H walkthroughs per month. CD has zero lost workday injuries for greater than one year. We have assessed ergonomics throughout one-third of the division and implemented corrections where needed. We have developed a total integrated safety management (ISM) program.

Goal: Manage and track computing equipment and instrumentation in a manner that supports the rest of the division's mission and works in concert with the laboratory's property tracking systems.

Assessment: The division collects property information and tracks property to support several mission functions that include the PREP equipment pool (support of the experimental program), administration of hardware and software maintenance contracts for the entire laboratory, tracking of all systems on the network and support for computer security. It is common to track the property through its entire life cycle...receiving, deployment to project, repair, off-site loans, on-site storage and retirement. The division is not required to track property to any DOE requirements when using the property information to support the above functions. The division is required to participate in property tracking processes for property that is stored at Site 38, property designated a sensitive item/capital equipment or property that is taken off-site.

The division takes actions required to track property information to the extent necessary to fulfill the division's mission. The division meets the requirements for property tracking as stated by the Laboratory/BSS/DOE (Sensitive items, etc.). The division propagates information on how to use property safely.

1) Improvements could be made with compliance of property passes. A reasonable procedure and on-line tools would improve compliance.
2) The division could take a lead role to develop on-line property transfer tools. This would greatly reduce the anguish each January as the inventory lists are sent out (in fact could eliminate the need for the paper lists).
3) The division could also develop tools that would allow the Laboratory's data repository to be augmented with the division's data.

4) The change management process for the division's property information is somewhat weak when changes are made in the field.